Distributed Self-Commanding Robotic Systems: A NASA CETDP Thinking Systems Thrust Proposal

Task Leaders

Anthony Barrett, JPL, anthony.barrett@jpl.nasa.gov, 818-393-5372, co-PI
Tara Estlin, JPL, tara.estlin@jpl.nasa.gov, 818-383-5393, co-PI

Product Description

This proposal involves developing technology for a system that continuously plans to control multiple spacecraft and/or rovers using collective mission goals instead of goals or command sequences for each robotic system. A population of self-commanding robots would autonomously coordinate itself to satisfy high level science and engineering goals in a changing partially-understood environment – making feasible the coordinated operation of tens or even a hundred spacecraft and rovers (such as for an interferometer, a magnetospheric constellation, or a Mars outpost).

This is a **new**, **push** task. (approximately 60% push and 40% pull)

Benefits

At the moment NASA is doing research on formation flying and networks of smart sensors for constellations, but no one is addressing the problem of autonomous constellation management. Without such autonomy managing a constellation of spacecraft either involves giving a sequence to one spacecraft and having it tele-operate the others or giving a separate sequence to each spacecraft. Unfortunately neither approach scales well with the constellation's population and complexity. While the first approach fails to scale due to bandwidth limitations, experience within the multi-agent research community has shown that significant numbers of unanticipated interactions between agents (like spacecraft) appear when people attempt to manually engineer sequences for more than 3 coordinated agents (Tambe 1997).

By reducing mission ops costs while enabling dynamic operation, this technology enhances missions involving multiple spatially-separated sensors needed for building large aperture virtual sensors, making in-situ measurements of macro phenomena, or exposing single sensors to risk without exposing the entire mission. Examples of such missions include:

- imaging extra-solar planets, observing sub-storms in the magnetosphere, and mapping geological formations on Mars (for the Space Science Enterprise);
- smart sensor networks on constellations of satellites to engage the public by directly responding to their science requests (for the Earth Science Enterprise); and
- teams of robotic partners for reconnaissance and work sharing (for Human Exploration and Development of Space).

Technical Approach

In general, autonomous spacecraft and rovers must balance long-term and short-term considerations. They must perform purposeful activities that ensure long-term science and engineering goals are achieved and ensure that they each maintain positive resource margins. This requires planning in advance to avoid a series of shortsighted decisions that can lead to failure. However, they must also respond in a timely fashion to a dynamic and unpredictable environment. In terms of high-level, goal-oriented activity, the robotic systems must modify their collective plans in the event of fortuitous events such as detecting scientific opportunities like a Martian hydro-thermal vent or a sub-storm onset in Earth's magnetosphere, and setbacks such as a spacecraft losing attitude control. For a single autonomous spacecraft, the software to satisfy these requirements can be partitioned into 4 components:

- a mission manager to generate high-level science goals from commands and detected opportunities,
- a planner/scheduler to turn goals into activities while reasoning about future expected situations,

- an executive/diagnostician to initiate and maintain activities while interpreting sensed events, and
- a conventional set of hardware proxies to interface with the spacecraft to implement an activity's primitive feedback loops.

While there are many approaches to coordinating a set of agents, the two most obvious either 1. treats them as a single master agent directing a set of slaves or 2. treats them as a set of competing peers. Actually, these two architectures determine a whole spectrum of approaches where the master gives its slaves progressively more autonomy. This progression manifests by giving the slaves more of the previously mentioned 4 components. In our work, we propose to start with the CASPER continuous planner (Chien et al. 1999) approach toward autonomy, implement both ends of the spectrum and then develop intermediate points. A continuous planner, like CASPER or IPEM (Ambros-Ingerson&Steel 1988), continuously extends and repairs a plan as activities execute and have unexpected results.

Uniformly Loose Coordination

Whether they are spacecraft, probes or rovers, coordinating multiple distributed agents introduces unique challenges for all four autonomy-supporting technologies. Issues arise concerning interfaces between agents, communication bandwidth, group command and control, and onboard capabilities. For example, consider a mission with a lander and a population of rovers for remote field geology. A certain level of communication capabilities will need to be assigned to each, possibly limiting the amount of information that can be shared between the rovers (and ground). The mission design will need to include a "chain of command" for the team of spacecraft/rovers, indicating which rovers are controlled directly from the ground, and which are controlled by other rovers or orbiting/landed spacecraft. The onboard capabilities also need consideration, including computing power and onboard data storage capacity. This will limit the level of autonomy each of the rovers can have. Finally, these issues apply to multiple spacecraft missions too – a constellation of orbiters has a ground station with a loosely coordinated population of satellites.

Many of these design issues are related, and all of them have an impact on possible automated planning and scheduling for the mission. The interfaces determine what activities can be planned for each robotic system. The amount of communication available will determine how much each can share its plan. The control scheme will also determine which spacecraft/rovers execute what activities in the plans. If one directs another, the "leader" will send activities from its plan to the "follower" for execution. Decisions on the onboard capabilities of a spacecraft/rover, however, will limit its independence. With little computing power, one spacecraft may be unable to plan and may only be able to execute commands. More power may allow it to plan and execute. Still more power may allow a spacecraft to plan for itself and others.

Assuming that each robotic system has enough computing power to plan for itself, we will develop several continuous goal-distribution techniques. In one approach, a leader uses a continuous planner (Chien 1999) to distribute goals to the followers who in turn use continuous planners to satisfy the distributed goals. Another approach has continuous planners on all of the followers, and the leader auctions off goals to the followers. While the first approach can always be crafted to provide superior solutions, it requires maintaining an abbreviated model of the whole constellation for the goal distribution planner. The second approach does not require this centralized model. Both approaches are amenable to robust execution. Given a centralized distribution planner, one rover/spacecraft discovering that it cannot satisfy a goal results in the central planner modifying its distribution plan to pass the goal to another robotic system. For the auction approach a failing rover can run its own auction to pass its objectionable goal off to another robotic system.

Loose and Tight Coordination

The easiest way to adapt autonomous spacecraft research to controlling tightly coordinated constellations involves treating the constellation as a single spacecraft. Here one spacecraft directly controls the others as if they were connected. The controlling "master" spacecraft performs all autonomy reasoning while the slaves only transmit sensor values to the master and forward control signals received from the master to their appropriate local devices. The executive/diagnostician starts actions and the master's reactive controller manages actions either locally or remotely through a slave. Adding slaves to our loose coordination model results in figure 1 where both leaders and followers can be masters (Barrett 1999).

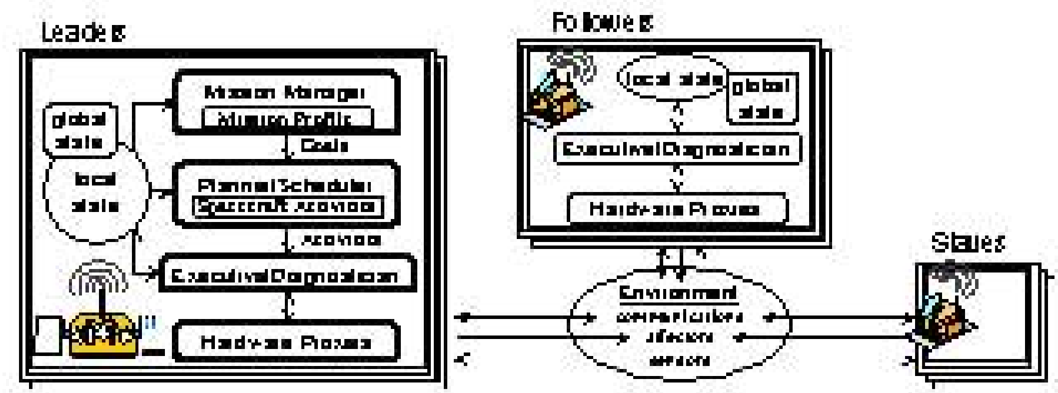

Figure 1. Software anatomy of leaders, followers, and slaves in an autonomous constellation

While the master/slave approach benefits from conceptual simplicity, it relies on an assumption that the master spacecraft's hardware proxies can continuously monitor the slaves' hardware, and this relies on high-bandwidth highly reliable communications. Since unintended results occur fairly rarely, one way to relax the bandwidth requirements involves putting hardware proxies on the slaves and only monitoring unexpected events. Unfortunately, experience within the multi-agent community shows that this approach disables the ability to monitor for unexpected events between spacecraft and leads to a host of coordination problems among the slaves (Tambe 1997). Upgrading the slaves to followers alleviates these problems. This upgrade results in a "teamwork" model of tight coordination. This model involves explicitly reasoning about local activities in relation to global "joint" activities and is currently a hot topic within the multi-agent community. We plan on adapting results from this research into our approach.

While our focus will initially fall on single missions with multiple leaders or a single leader with multiple slaves possibly upgraded to followers, our ultimate goal is to provide an infrastructure where an evolving mission can have leaders, followers and slaves. Such an infrastructure would support a robotic colony where elements degrade and are periodically re-supplied. Another added complexity involves the observation that missions tend to have multiple PIs developing different sensors. For a multi-platform mission, these distributed sensors might have different coordination requirements. For instance, a constellation might implement an optical interferometer where each spacecraft also has a plasma physics module. While the interferometer needs tight coordination, the other modules only need loose coordination. In such a case all spacecraft function as leaders for the plasma physics experiments, but only one functions as a leader for the interferometer.

Current State of The Art

The concept of distributed self-commanding robotic systems is not new, and we can characterize our approach in terms of combining ideas from several systems described within the established multi-agent literature. INTERRAP (Müller 1996), LEMMING (Ohko et al. 1995), and GRAMMPS (Brumitt&Stentz 1998) each address problems involving loose coordination between a population of mobile robots, but none of these systems address problems involving tight coordination among a number of agents where individual agents can fail. STEAM (Tambe 1997) and TPOT-RL (Stone 1998) are two systems that address teams of *tightly coordinated agents that can fail, but they primarily focus on executing plans and responding to failures. With respect to figure 1, these systems focus on coordinating a set of followers. They do not address autonomously building or repairing plans. Finally, none of these systems plan with temporal and resource constraints. While CASPER (Chien 1999), IxTeT (Laborie 1995), and RAX (Muscettola et al. 1998) address temporal and resource constraints, they do not address robust cooperation.*

While all of the above systems satisfy subsets of our problem, none of them will autonomously control multiple robotic systems in the presence of temporal and resource constraints with loose and tight coordination requirements.

Status and Milestones

While this is a new task, the PIs have been involved in related tasks for the past year. One collaboration resulted in a study of 3 coordination approaches where a lead spacecraft progressively downloads the planning burden onto its followers (Rabideau et al. 1999). Figure 2 illustrates these approaches with the leader is on the lander, but the leader could run on either the ground or one of the rovers. In each case the collection of planning processes interacts to maximize the number of observations made by the rovers while minimizing the distance traveled. We experimented with these 3 points using the ASPEN planner (Fukunaga et al. 1997)– a non-continuous predecessor of CASPER.

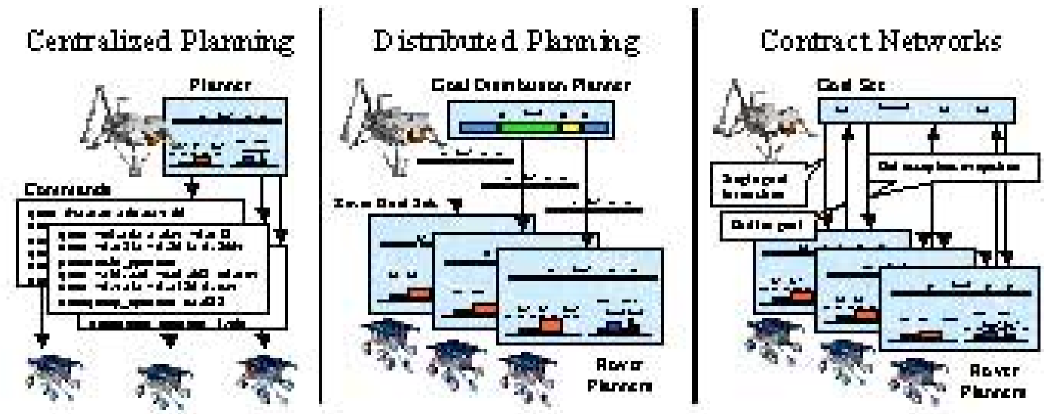

Figure 2. Coordination Architectures for a set of Rovers

In the centralized planning case, the lander plans to command the slave rovers with sequences. While this is the simplest conceptual extension to ASPEN, applying it to CASPER will suffer from the amount of communication required to continuously micro-manage the rovers. Using a goal distribution planner for distributed planning reduces the bandwidth requirement. In this case the lander has abbreviated models of the rovers and plans to the point where it can determine the appropriate rover for each observation. After transmitting the observation request, the receiving rover fills in the details in its local plan. Giving the rovers even more autonomy results in replacing the distribution planner with a contract network. Instead of telling which rover to satisfy a goal, this approach advertises the goal and lets the rovers bid for it in accordance with how well it can insert the goal into its current plan.

FY 2000 Milestones:

- Demonstrate goal-level commanding of a set of 3 to 5 simulated spacecraft or rovers where each robotic system's activities do not directly interact with another's (loose coordination).

FY 2001 Milestones:

- Demonstrate the above with larger numbers of simulated robotic systems.

- Demonstrate scenarios where 2 to 5 robotic systems can tightly coordinate to perform small joint activities (loose and tight coordination).

FY 2002 Milestones:

- Demonstrate on real rovers and detailed simulations of distributed spacecraft.
- Expand demonstration to allow larger joint activities.

Customer Relevance

This proposal has been discussed with two JPL project geologists, Dr. Steve Saunders (Mars-01 Project Scientist) and Dr. Ashley Davies (NIMS/Galileo Project Scientist) – support letter attached. They both agree that this work is important for future rover missions to provide more autonomous capabilities for teams of rovers/spacecraft and for furthering planetary science experiments. From a more constellation oriented perspective, Dr. Michael Rilee (GSFC Plasma Physicist, Solar-Terrestrial Probe Line Science Application Team, Remote Exploration & Experimentation Project) agrees that this work is important for future multi-platform missions proposed for the next generation of solar-terrestrial probes – support letter attached.

This proposal has also been endorsed by Samad Hayati (Manager of the Robotics and Mars Exploration Technology Program) and who believes this technology will play a key role in deploying distributed robotic systems to Mars and in other relevant future missions – support letter attached. We have also discussed this proposal with Tom Starbird (Lead, Execution & Planning Domain, Mission Data Systems Project) – support letter attached. He agrees that this work aligns well with the Mission Data System architecture and is important for illuminating the issues involved in extending the architecture to multi-platform missions.

Technical References

J. Ambrose-Ingerson, S. Steel, “Integrating Planning, Execution and Monitoring,” Proceedings of the Seventh National Conference on Artificial Intelligence, Saint Paul, Minnesota, August 1988.

A. Barrett, "Autonomy Architectures for a Constellation of Spacecraft," International Symposium on Artificial Intelligence Robotics and Automation in Space (iSAIRAS), Noordwijk, The Netherlands, June 1999.

B. Brumitt,, A. Stentz, "GRAMMPS: A Generalized Mission Planner for Multiple Mobile Robots", Proceedings of the IEEE International Conference on Robotics and Automation, May 1998.

S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Using Iterative Repair to Increase the Responsiveness of Planning and Scheduling for Autonomous Spacecraft," IJCAI99 Workshop on Scheduling and Planning meet Real-time Monitoring in a Dynamic and Uncertain World, Stockholm, Sweden, August 1999.

T. Estlin, A. Gray. T. Mann, G. Rabideau, R. Castaño, S. Chien, and E. Mjolsness, “An Integrated System for Multi-Rover Scientific Exploration,” Proceedings of the Sixteenth National Conference on Artificial Intelligence, Orlando, FL, July 1999.

T. Estlin, S. Hayati, A. Jain, J. Yen, G. Rabideau, R. Castano, R. Petras, S. Peters, D. Decoste, E. Tunstel, S. Chien, E. Mjolsness, R. Steele, D. Mutz, A. Gray, T. Mann, "An Integrated Architecture for Cooperating Rovers," International Symposium on Artificial Intelligence Robotics and Automation in Space (ISAIRAS), Noordwijk, The Netherlands, June 1999.

A. Fukunaga, G. Rabideau, S. Chien, D. Yan, "Towards an Application Framework for Automated Planning and Scheduling," Proceedings of the 1997 International Symposium on Artificial Intelligence, Robotics and Automation for Space, Tokyo, Japan, July 1997.

P. Laborie, M. Ghallab, “Planning with Sharable Resource Constraints,” Proceedings of the International Joint Conference on Artificial Intelligence, Montreal, Canada, August 1995.

J. Müller, *The Design of Intelligent Agents, A Layered Approach*. Lecture Notes in Artificial Intell-igence, Springer-Verlag, 1996.

N. Muscettola, P. Nayak, B. Pell, and B. Williams, “Remote Agent: To Boldly Go Where No AI System Has Gone Before,” Artificial Intelligence 103(1-2):5-48, August 1998.

G. Rabideau, T. Estlin, S. Chien, A. Barrett, "A Comparison of Coordinated Planning Methods for Cooperating Rovers," AIAA 1999 Space Technology Conference, Albuquerque, NM, September 1999.

G. Rabideau, S. Chien, T. Estlin, A. Barrett, "Working Together: Centralized Command Sequence Generation for Cooperating Rovers," Proceedings of the IEEE Aerospace Conference (IAC), Aspen, CO, March 1999.

P. Stone. Layered Learning in Multi-Agent Systems. PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, December 1998. Available as technical report CMU-CS-98-187

M. Tambe, “Towards Flexible Teamwork,” *Journal of Artificial Intelligence Research*, 7:83-124.

June 24, 1999

Thinking Systems Thrust Area, Cross Enterprise Technology Development Program

Subject: **Distributed Self-Commanding Robotic Systems (DSCRS) Task**

Reference: **Letter of Support**

The purpose of this letter is to indicate our support for the work in development and deployment of Distributed Planning and Scheduling Technology for commanding constellations of spacecraft and populations of rovers funded by the Thinking Systems Thrust Area of the CETDP and lead by Dr. Anthony Barrett and Dr. Tara Estlin.

The aim of the Solar-Terrestrial Probe Line Science Application Team of NASA's Remote Exploration and Experimentation (REE) project is to enable classes of missions that are currently impossible due to lack of communication and computational resources. Multispacecraft missions are a key element in the Space Science Enterprise's Sun-Earth Connections Theme. Yet such missions will place extraordinary demands on the communication and control infrastructures required to operate these systems. Coordinated, intelligent control of these instruments is also critical to the success of these missions.

Cluster II is a four spacecraft observatory will be launched in 2000 and will dramatically improve our knowledge of the Sun-Earth space environment. Two Solar-Terrestrial Probe Line observatories, currently envisioned at four spacecraft each, *Magnetospheric Multiscale* (2007) and *Global Electrodynamic Connections* (2008), will target specific regions of the Earth's magnetosphere and ionosphere. These missions will have sensitive magnetometers and plasma wave detectors that provide information that could be used to autonomously redirect mission science activities. Currently this capability is extremely limited. Even the simple coordination of data collection, selection, and transmission can greatly enhance science return.

The work proposed by Barrett and Estlin may help reduce command and communication costs, thus enabling certain deep space missions or those missions with many spacecraft elements. A fundamental requirement for missions such as *Magnetospheric Constellation* (2010) is that a constellation of 100 spacecraft must be no more difficult to operate than one spacecraft today. The proposed effort is definitely applicable to advanced mission concepts such as the *Solar Imaging Radio Array*. *SIRA* must be placed far from the Earth and has a raw data production rate measured in hundreds of gigabits per day. Tight coordination amongst the array elements is required for radio imaging to be accomplished.

Because of the alignment of the DSCRS task with our objective to autonomously extract and downlink information relevant to mission goals, there is a strong reason to expect that this work will significantly contribute to the next generation of Solar-Terrestrial Probes. We strongly support funding for the DSCRS task.

Michael L. Rilee, Raytheon, GSFC, 301-286-4743, Michael.L.Rilee.1@gsfc.nasa.gov.
For the Solar-Terrestrial Probe Line Science Application Team of NASA's Remote Exploration and Experimentation Project, S.A. Curtis, NASA, Head, Code 695, GSFC.

June 23, 1999

To: Thinking Systems Thrust Area, Cross Enterprise Technology Development Program

From: T. Starbird,
Lead, Execution & Planning Domain, Mission Data Systems Project

Subject: **Distributed Self-Commanding Robotic Systems (DSCRS) Task**

The purpose of this letter is to indicate my support for the work in development and deployment of Distributed Planning and Scheduling Technology for commanding constellations of spacecraft and populations of rovers funded by the Thinking Systems Thrust Area of the CETDP and lead by Dr. Anthony Barrett and Dr. Tara Estlin.

In my role as lead of the Execution & Planning Domain of the Mission Data Systems (MDS) Project, I am responsible for developing the Goal-based Control Architecture (called the GAM architecture). Current plans are for the basic GAM architecture to rely solely on hard-coded GAMs for automation. But the design of the GAM architecture will be one that allows plugging a continuous planner (like CASPER) onboard a spacecraft, thereby enabling its future use as the technology is accepted by flight projects. This technology insertion path is documented in the Advanced Planning & Sequencing Technology Development plan that I wrote (January 1999). The effort in the DSCRS task to coordinate multiple continuous planners distributed across multiple platforms will define an extra step in this path that will generalize the MDS to apply to constellation class missions.

The current MDS customer missions involve a single spacecraft each, but the architecture is consciously and specifically designed to handle multi-platformed missions as well. Technology from the DSCRS task will contribute significantly towards illuminating the issues involved, and extending the MDS architecture to multi-platform missions; I support funding for the DSCRS task.

June 22, 1999

To: **Evaluation Committee,**
Thinking Systems Thrust Area,
Cross Enterprise Technology Development Program

Subject: **Support for Distributed Self-Commanding Robotic Systems Proposal**

The purpose of this letter is to indicate my support for developing automated planning and scheduling technology to control multiple rovers and/or spacecraft, as described in the Distributed Self-Commanding Robotic Systems Proposal, led by Dr. Tara Estlin and Dr. Anthony Barrett at JPL. This task has been submitted to the Thinking Systems Thrust Area of the CETDP.

I have worked with Dr. Tara Estlin as part of the Advanced Cooperating Rovers Task (funded through the Advanced Concepts Program), which has provided preliminary funding to develop a cooperative software architecture for autonomous rovers. Tara has been the enabling force in allowing the ACR team to take advantage of AI planning and scheduling technology in general and has led the effort for this task to develop and integrate a distributed version of the CASPER real-time planning system to coordinate and control multiple rovers. The use of this system has contributed significantly to the ACR task by providing abilities for automatic command-sequence generation, resource and constraint modeling, goal distribution and dynamic re-planning.

In my role as manager of the Robotics and Mars Exploration Technology Program, I foresee this technology playing a critical role in future missions that involve multiple robotic systems. Distributed planning technology allows a team of rovers to autonomously coordinate itself to satisfy science and engineering goals in a changing, partially-understood environment. This technology has a number of important benefits for any distributed robotic system. First, it greatly reduces the cost of command-sequence generation and validation. Currently this process is a very expensive and labor-intensive for single rover operations and will be much more complicated and costly for large numbers of rovers. Second, onboard autonomous systems will be more critical for teams of rovers which must constantly be communicating and sharing information in order to accomplish mission goals. Third, planning technology allows a team of rovers to more responsive to unexpected changes in the environment which allows more robust behavior and enables significantly more opportunities for science.

In closing, I believe that automated planning and scheduling technology will play a key role in deploying distributed robotic systems to Mars and in other relevant future missions. Thus, I strongly support funding the Distributed Self-Commanding Robotics Systems Proposal.

I am willing to provide funding towards the testing of these technologies on realistic rovers at JPL.

Sincerely,

Samad Hayati

Samad Hayati
Manager, Robotics and Mars Exploration Technology
JPL Representative, Surface Systems Thrust Area, CETDP

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109-8099
(818) 354-4321

JPL

To:	The Evaluation Committee, Thinking Systems Thrust Area, Cross Enterprise Technology Development Program
From:	Ashley G. Davies ms 183-601 Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109.
Subject:	Support for "Distributed Self-Commanding Robotic Systems" Proposal
Date:	24 June 1999

To whom it may concern:

This is a letter in support of Dr. Tara Estlin's proposal "Distributed Self-Commanding Robotic Systems". I have worked with Dr. Estlin as part of the Advanced Cooperating Rovers Task (funded through the Advanced Concepts Program), which provided preliminary funding to develop a cooperative software architecture for autonomous rovers. Tara has been the main force in allowing the Advanced Co-operating Rovers team to take advantage of artificial intelligence planning and scheduling technology. The value of the this system to JPL has been recognised by the bestowal of a NOVA Team award to the Multi-Rover Integrated Science Understanding System (MISUS) Team, of which Tara is a lead element.

As a geologist, I feel that the work Tara is doing, and is proposing to do, is of great value to JPL in carrying out it's goal of remotely exploring the Solar System. From the perspective of wanting to explore geological systems on Mars, for example, networks of rovers will be used to pave the way for, and then augment, a human presence on Mars. In order for the rovers to perform with the highest efficiency to maximise science return, the rover network must be controlled at two levels: from the point of view of the individual instrument or rover; and from the point of view of the overall mission. This proposal builds on the work Tara has already done concerning the latter: the control of self-commanding, autonomous multiple rover and/or spacecraft to fulfill mission goals.

Such a system is vastly more time efficient than current command sequence and validation processes. A network of rovers can economically and efficiently carry out mission tasks with science returns way beyond what individual rovers can accomplish, as long as the multiple rovers are controlled to augment each others science return. In this respect, the work Dr. Estlin is proposing is of great value, and has my strongest support.

Yours faithfully,

Ashley G. Davies, Ph.D.
Scientist, Galileo-NIMS
Co-operating Rovers for Remote Field Geology Project